# A Hybrid Gradient Method to Designing Bayesian Experiments for Implicit Models

Jiaxin Zhang Computer Science and Mathematics Division Oak Ridge National Laboratory, Oak Ridge, TN 37830 zhangj@ornl.gov

Sirui Bi Computational Sciences and Engineering Division Oak Ridge National Laboratory, Oak Ridge, TN 37830 bis1@ornl.gov

Guannan Zhang Computer Science and Mathematics Division Oak Ridge National Laboratory, Oak Ridge, TN 37830 zhangg@ornl.gov

# Abstract

Bayesian experimental design (BED) aims at designing an experiment to maximize the information gathering from the collected data. The optimal design is usually achieved by maximizing the mutual information (MI) between the data and the model parameters. When the analytical expression of the MI is unavailable, e.g., having implicit models with intractable data distributions, a neural network-based lower bound of the MI was recently proposed and a gradient ascent method was used to maximize the lower bound [1]. However, the approach in [1] requires a pathwise sampling path to compute the gradient of the MI lower bound with respect to the design variables, and such a pathwise sampling path is usually inaccessible for implicit models. In this work, we propose a hybrid gradient approach that leverages recent advances in variational MI estimator and evolution strategies (ES) combined with black-box stochastic gradient ascent (SGA) to maximize the MI lower bound. This allows the design process to be achieved through a unified scalable procedure for implicit models without sampling path gradients. Several experiments demonstrate that our approach significantly improves the scalability of BED for implicit models in high-dimensional design space.

# 1 Introduction

Experimental design plays an essential role in all scientific disciplines. Our ultimate goal is to determine designs that maximize the information gathered through the experiments so that improve our understanding on model comparison or parameter estimations. A broadly used approach is Bayesian experimental design (BED) [2] that aims to find an optimal design  $\boldsymbol{\xi}^*$  to maximize a utility function  $I(\boldsymbol{\xi})$ , which is typically defined by the mutual information (MI) between data and model parameters. Typically, the BED framework begins with a Bayesian model of the experimental process, including a prior distribution  $p(\boldsymbol{\theta})$  and a likelihood  $p(\boldsymbol{y}|\boldsymbol{\theta}, \boldsymbol{\xi})$ . The information gained about  $\boldsymbol{\theta}$  from running the experiment with design  $\boldsymbol{\xi}$  and observed outcome  $\boldsymbol{y}$  can be interpreted by the reduction in

Third Workshop on Machine Learning and the Physical Sciences (NeurIPS 2020), Vancouver, Canada.

entropy from the prior to posterior

$$IG(y, \boldsymbol{\xi}) = \mathcal{Q}[p(\boldsymbol{\theta})] - \mathcal{Q}[p(\boldsymbol{\theta}|y, \boldsymbol{\xi})].$$
(1)

To define a metric to quantify the utility of the design  $\boldsymbol{\xi}$  before running experiments, an expected information gain (EIG),  $I(\boldsymbol{\xi})$  is often used:

$$I(\boldsymbol{\xi}) = \mathbb{E}_{p(\boldsymbol{y}|\boldsymbol{\xi})}[\mathcal{Q}[p(\boldsymbol{\theta})] - \mathcal{Q}[p(\boldsymbol{\theta}|\boldsymbol{y},\boldsymbol{\xi})]].$$
(2)

Eq. (2) can also be interpreted as a mutual information (MI) between  $\theta$  and y with a specified  $\xi$ ,

$$I_{\rm MI}(\boldsymbol{\xi}) = \mathbb{E}_{p(\boldsymbol{\theta})p(\boldsymbol{y}|\boldsymbol{\theta},\boldsymbol{\xi})} \left[ \log \frac{p(\boldsymbol{y}|\boldsymbol{\theta},\boldsymbol{\xi})}{p(\boldsymbol{y}|\boldsymbol{\xi})} \right]$$
(3)

The Bayesian optimal design is therefore defined as

$$\boldsymbol{\xi}^* = \operatorname*{arg\,max}_{\boldsymbol{\xi} \in \boldsymbol{\Xi}} I_{\mathrm{MI}}(\boldsymbol{\xi}) \tag{4}$$

where  $\Xi$  is the feasible design domain. The most challenging task in BED framework is how to efficiently and accurately estimate  $I_{\text{MI}}(\boldsymbol{\xi})$  in Eq. (3) and optimize  $I_{\text{MI}}(\boldsymbol{\xi})$  via Eq. (4) to obtain the optimal design  $\boldsymbol{\xi}^*$ .

Most of the existing BED studies focus on the *explicit models* [2, 3, 4, 5] in which the likelihood is analytically known, but in natural and physical science, a more common scenario is the *implicit models* [6, 7], in which the likelihood is intractable but sampling is possible. In other words, the implicit model is specified based on a stochastic data generating simulator and typically has no access to the analytical form and the gradients of the joint density  $p(\theta, y|\xi)$  and marginal density  $p(y|\xi)$ . The resulting BED scheme shares a two-stage feature: build a pointwise estimator of  $I(\xi)$  and then feed this "black-box" estimator to a separate outer-level optimizer such as Bayesian optimization to find the optimal design  $\xi^*$ . This scheme substantially increases the overall computational cost and is challenging in scaling the BED to a high dimensional design space.

Recent studies [1, 5, 8] alleviate the challenges by using stochastic gradient-based approaches but they rely on the models with tractable likelihood functions or assume the gradients can be reasonably approximated by pathwise gradient estimators with sampling path, unlike the scope of our paper that focuses on the BED for implicit models without gradients. We develop a general scalable framework that can jointly optimize a unified objective with respect to both the variational MI bound and the design using a hybrid gradient ascent approach. The key contributions are summarized as follows:

- We propose a hybrid gradient approach that leverages recent advances in variational MI estimator and guided evolution strategies (ES) to maximize the MI lower bound;
- We incorporate a smoothed estimator of the MI lower bound via clipped density ratios to reduce the variance in the MI bound estimator and the estimation of the posterior samples;
- We demonstrate the superior performance of our proposed approach on a toy noisy linear problem and a real quantum control problem, specifically in high dimensional design spaces.

# 2 Hybrid Gradient Method for BED

Here, we aim to find the optimal design  $\xi^*$  by maximizing a MI lower bound rather than exactly estimating MI with high accuracy. Estimating and optimizing MI is core to many machine learning research but it has been a challenge to bounding MI in high dimensions.

**Mutual information estimators** Belghazi et al. [9] proposed to estimate the MI using gradient descent over neural networks and argued that the lower bound can be tightened by optimizing the neural network parameters. The MI estimator is typically named by MINE-*f* or *f*-GAN KL [10]

$$I_{\text{MINE}}(\boldsymbol{\xi}, \boldsymbol{\psi}) = \mathbb{E}_{p(\boldsymbol{\theta}, \boldsymbol{y} | \boldsymbol{\xi})}[\mathcal{T}_{\boldsymbol{\psi}}(\boldsymbol{\theta}, \boldsymbol{y})] - \log \mathbb{E}_{p(\boldsymbol{\theta})p(\boldsymbol{y} | \boldsymbol{\xi})}[e^{\mathcal{T}_{\boldsymbol{\psi}}(\boldsymbol{\theta}, \boldsymbol{y})}]$$
(5)

where  $\mathcal{T}_{\psi}(\theta, y)$  is a neural network that is parametrized by  $\psi$  with model parameters  $\theta$  and data y as inputs. Incorporating neural network parameters  $\psi$  with design parameters  $\xi$ , the BED problem can be formulated by maximizing the overall objective

$$\boldsymbol{\xi}^* = \arg\max_{\boldsymbol{\xi}} \max_{\boldsymbol{\psi}} \left\{ I_{\text{MINE}}(\boldsymbol{\xi}, \boldsymbol{\psi}) \right\}.$$
(6)

The optimal design  $\boldsymbol{\xi}^*$  is obtained by maximizing the MI estimator in Eq. (5) through a joint gradientbased algorithm or a separate gradient-free update scheme of  $\boldsymbol{\xi}$  and  $\boldsymbol{\psi}$ . The effectiveness for the MI estimation and optimization therefore becomes very important to the BED problem. Unfortunately,  $I_{\text{MINE}}$  exhibits a high variance that could grows exponentially with the ground truth MI and leads to poor bias-variance trade-offs in practice [11, 12]. To address the high-variance issue in the  $I_{\text{MINE}}$ estimator, we propose to use a smoothed MI lower-bound estimator  $I_{\text{SMILE}}$  [12] with hyperparameter  $\tau$  that clips the density ratios when estimating the partition function:

$$I_{\text{SMILE}}(\boldsymbol{\xi}, \boldsymbol{\psi}) = \mathbb{E}_{p(\boldsymbol{\theta}, \boldsymbol{y} | \boldsymbol{\xi})}[\mathcal{T}_{\boldsymbol{\psi}}(\boldsymbol{\theta}, \boldsymbol{y})] - \log \mathbb{E}_{p(\boldsymbol{\theta})p(\boldsymbol{y} | \boldsymbol{\xi})}[\operatorname{clip}(e^{\mathcal{T}_{\boldsymbol{\psi}}(\boldsymbol{\theta}, \boldsymbol{y})}, e^{-\tau}, e^{\tau})]$$
(7)

where clip function  $\operatorname{clip}(\mathbf{u}, \mathbf{v}, \mathbf{w}) = \max(\min(\mathbf{u}, \mathbf{w}), \mathbf{v})$ . The choice of  $\tau$  affects the bias-variance trade-off: when  $\tau \to \infty$ ,  $I_{\text{SMILE}}$  converges to  $I_{\text{MINE}}$ ; with a smaller  $\tau$ , the variance is reduced at the cost of increasing bias [12]. The improved MI estimation via variance reduction techniques is benefit to the optimization process in Eq. (7) and thus leads to a robust final optimal design  $\boldsymbol{\xi}^*$ .

**Evolutionary strategies** When the gradient of the MI lower bound is inaccessible, a popular zero-order approach for estimating the gradient is Gaussian Smoothing (GS) (or called Evolution Strategies [13]). The smoothed loss is defined by

$$f_{\sigma}(\boldsymbol{\xi}) = \mathbb{E}_{\boldsymbol{\epsilon} \sim \mathcal{N}(0, \mathbf{I}_d)} \left[ f(\boldsymbol{\xi} + \sigma \boldsymbol{u}) \right] = (2\pi)^{-n/2} \int_{\mathbb{R}^n} f(\boldsymbol{\xi} + \sigma \boldsymbol{\epsilon}) e^{-\|\boldsymbol{\epsilon}\|^2/2} d\boldsymbol{\epsilon}$$
(8)

where  $\mathcal{N}(0, \mathbf{I}_d)$  is the *d*-dimensional standard Gaussian distribution, and  $\sigma > 0$  is the smoothing radius. The standard GS represents the  $\nabla f_{\sigma}(\boldsymbol{\xi})$  as an *d*-dimensional integral and estimate it by drawing *M* random samples  $\{\boldsymbol{\epsilon}_i\}_{i=1}^M$  from  $\mathcal{N}(0, \mathbf{I}_d)$ , i.e.,

$$\nabla f_{\sigma}(\boldsymbol{\xi}) = \frac{1}{\sigma} \mathbb{E}_{\boldsymbol{\epsilon} \sim \mathcal{N}(0, \mathbf{I}_d)} \left[ f(\boldsymbol{\xi} + \sigma \boldsymbol{\epsilon}) \, \boldsymbol{\epsilon} \right] \approx \frac{1}{M\sigma} \sum_{i=1}^M f(\boldsymbol{\xi} + \sigma \boldsymbol{\epsilon}_i) \boldsymbol{\epsilon}_i. \tag{9}$$

The MC estimator in Eq. (9) is usually used as an unbiased estimator of the local gradient  $\nabla F(\mathbf{x})$  by exploiting the fact that  $\lim_{\sigma \to 0} \nabla F_{\sigma}(\mathbf{x}) = \nabla F(\mathbf{x})$ , where the smoothing radius  $\sigma$  is often set to a small value. However, the traditional GS tends to a high variance for high dimensional space.

Several advances in evolution strategies (ES) [14, 15, 16, 17] addressed these issues using variance reduction and dimension reduction strategies. More recently, a Guided ES method [18] is proposed by optimally using surrogate gradient directional along with a random search. Specifically, Guided ES generates a subspace by keeping track of the previous k surrogate gradients during optimization, and leverages this prior information by changing the distribution of  $\epsilon_i$  in Eq. (9) to  $\mathcal{N}(0, \Sigma)$  with  $\Sigma = (\alpha/n) \cdot \mathbf{I}_n + (1-\alpha)/k \cdot UU^T$  where k and n are the subspace and parameter dimensions respectively, U denotes an  $n \times k$  orthonormal basis for the subspace, and  $\alpha$  is a hyperparameter that trades off variance between the subspace and full parameter space. The improved search distribution allows a low-variance estimate of the descent direction  $\nabla f_{\sigma}^{c}(\boldsymbol{\xi})$ .

Stochastic approximate gradient ascent (SAGA) method We propose a hybrid gradient method by maximizing  $I_{\text{SMILE}}(\boldsymbol{\xi}, \boldsymbol{\psi})$  by gradient ascent on NN parameters  $\boldsymbol{\psi}$  combined with an approximate gradient  $\nabla f_{\sigma}^{G}(\boldsymbol{\xi})$  using Guided ES on design parameters  $\boldsymbol{\xi}$ . Starting from an initial design  $\boldsymbol{\xi}_{0}$  and NN parameters  $\boldsymbol{\psi}_{0}$ , the approximate gradient of the MI lower bound  $\nabla_{\boldsymbol{\xi}} I_{\text{SMILE}}(\boldsymbol{\xi}, \boldsymbol{\psi})$  at  $\boldsymbol{\xi}_{0}$  can be estimated by the Guided ES method. With the gradient in hand, we can maximize the MI lower bound with respect to both the NN parameters  $\boldsymbol{\psi}$  and the design parameters  $\boldsymbol{\xi}$  jointly by gradient ascent optimizer, which is given by Algorithm 1. The proposed method is named by SAGABED, which is a critical contribution to scale the Bayesian experimental design to high dimensional setting such that we can overcome the grand challenge in gradient-free methods, e.g., Bayesian optimization.

In the following, we discuss some important features of the SAGABED method, specifically for high-dimensional design problems: (1) *Unified framework vs. two-stage framework*: Without the requirement of pathwise gradients for implicit models, we utilize the stochastic approximate gradients and construct a unified framework that allows the design process to be performed by a simultaneous optimization with respect to both the variational and design parameters. The existing two-stage framework that builds a pointwise MI estimator before feeding this estimator to an outer-level optimizer is often computationally intensive; (2) *Scalability, portability, and parallelization*: we propose a stochastic approximate gradient ascent procedure that naturally avoids the scalability issue in gradient-free methods. The proposed framework can be easily incorporated with other MI estimators and implicit models because we only need the function value to approximate the gradient

based on the guided ES algorithm; (3) Robust estimation with a low variance: The smoothed MI lower bound used here allows us to perform a robust MI estimation and optimization of experimental design. The resulting low variance of the optimal design and posterior samples enable a more accurate estimate of the model parameters.

### Algorithm 1: The SAGABED algorithm

1: **Require**: neural network architectures, learning rates  $\ell_{\psi}$  and  $\ell_{\xi}$ ,  $\tau$  in  $I_{\text{SMILE}}$ , total prior samples n, total iterations T, implicit model  $\mathcal{M}$ 2: Process: 3: Initialize a design  $\boldsymbol{\xi}_0$  by random sampling 4: Initialize neural network parameter  $\psi_0$ 5: for t = 0 : T - 1 do Draw n samples from the prior distribution of the model parameters  $\theta$ :  $\theta^{(1)}, ..., \theta^{(n)} \sim p(\theta)$ 6: Compute the data samples  $y^{(i)}$ , i = 1, ..., n using the current design  $\xi_t$  and a implicit model  $\mathcal{M}$ 7: Evaluate the smoothed MI lower bound  $I_{\rm SMILE}$  at the current design  $\boldsymbol{\xi}_t$  and network parameters  $\boldsymbol{\psi}_t$ 8: 9: Compute the approximate gradient estimator  $\nabla_{\boldsymbol{\xi}} I_{\text{SMILE}}(\boldsymbol{\xi}_t, \boldsymbol{\psi}_t)$  using the GES algorithm 10: Evaluate the gradient of the  $I_{\text{SMILE}}$  with respect to the network parameters  $\nabla_{\psi} I_{\text{SMILE}}(\boldsymbol{\xi}, \boldsymbol{\psi})$ 11: Update design  $\boldsymbol{\xi}_t$  via gradient ascent:  $\boldsymbol{\xi}_{t+1} = \boldsymbol{\xi}_t + \ell_{\boldsymbol{\xi}} \nabla_{\boldsymbol{\xi}} I_{\text{SMILE}}(\boldsymbol{\xi}_t, \boldsymbol{\psi}_t)$ Update neural network parameters  $\psi_t$  via gradient ascent: 12:  $\boldsymbol{\psi}_{t+1} = \boldsymbol{\psi}_t + \ell_{\psi} \nabla_{\boldsymbol{\psi}} I_{\text{SMILE}}(\boldsymbol{\xi}_t, \boldsymbol{\psi}_t)$ 

```
13: end for
```

After determining the optimal design  $\boldsymbol{\xi}^*$  by maximizing the MI lower bound, we can obtain an estimate of the posterior  $p(\theta|y, \xi^*)$  given the learned neural network  $\mathcal{T}_{\psi^*}(\theta, y)$  and prior distribution

$$p(\boldsymbol{\theta}|\boldsymbol{y},\boldsymbol{\xi}) = \operatorname{clip}(e^{\mathcal{T}_{\boldsymbol{\psi}}(\boldsymbol{\theta},\boldsymbol{y})-1}, e^{-\tau}, e^{\tau})p(\boldsymbol{\theta})$$
(10)

The relationship in Eq. (10) allows to easily generate posterior samples  $\theta_i \sim p(\theta|\mathbf{y}, \boldsymbol{\xi}^*)$  using MCMC algorithm since the posterior density can be quickly evaluated via Eq. (10).

#### 3 **Experiments**

Noise linear regression We first demonstrate our proposed method using a classical noisy linear model [1], which is given by  $\boldsymbol{y} = \theta_1 \mathbf{1} + \theta_2 \boldsymbol{d} + \boldsymbol{\epsilon} + \boldsymbol{\nu}$ , where  $\boldsymbol{y}$  is a response variable,  $\boldsymbol{\theta} = [\boldsymbol{\theta}_1, \theta_2]^T$  are model parameters,  $\boldsymbol{\epsilon} \sim \mathcal{N}(0, 1)$  and  $\boldsymbol{\nu} \sim \Gamma(2, 2)$  are noise terms. Our target is to make D measurements to estimate  $\theta$  by constructing a design vector  $\boldsymbol{d} = [d_1, ..., d_D]^T$  which consists of individual experimental designs. Four cases D=1, 10, 50 and 100 are investigated in this example. We randomly initialize a design  $d \in [-10, 10]$  and sample 10,000 parameters from a prior distribution  $p(\theta) = \mathcal{N}(0, 3^2)$  and we use one layer of 100 and 150 hidden neurons for D = 1 and D = 10respectively. For high dimensional cases D = 50 and D = 100, we use 5-layered network with 50 hidden neurons for each layer. These NN architecture settings follow up the guidance in [1].



Figure 1: MI lower bound as a function of NN training epochs for D=1, 10, 50 and 100.

Figure 1 shows a comparison between the proposed SAGABED method (combined SMILE with Guided ES) and two baselines (MINE with SGA and BO). When D=1, three methods perform similarly but our method outperforms the other two baselines on the high dimensional cases. The SGA method is close to our method but shows a unstable training with large variance, and is probably trapped into a local minimal when D=10. The performance of the MINE with BO method drops significantly as the dimension increases. The proposed SAGABED method also demonstrates superior performance

Table 1: Estimating mean and standard deviation of the posterior samples of the model parameters  $\theta$  using optimal designs  $d^*$  and real data observation  $y^*$  (use  $\theta_{true} = [1,4]$  to generate  $y^*$ )

Method	D=10		D=50		D=100	
	$\hat{ heta_1}$	$\hat{ heta_2}$	$\hat{ heta_1}$	$\hat{ heta_2}$	$\hat{ heta_1}$	$\hat{ heta_2}$
MINE + SGA	$0.51 {\pm} 0.44$	$2.99 {\pm} 0.67$	$1.20 {\pm} 0.18$	$3.79 {\pm} 0.23$	$0.97 {\pm} 0.05$	$4.04 {\pm} 0.04$
MINE + BO	$1.22 {\pm} 0.58$	$4.93 {\pm} 0.91$	$0.71 {\pm} 0.25$	$3.66 {\pm} 0.40$	$1.35 {\pm} 0.11$	$4.79 {\pm} 0.26$
SMILE + GES	$0.83{\pm}0.56$	$4.69{\pm}0.58$	$1.11 \pm 0.13$	$4.25{\pm}0.19$	$1.02{\pm}0.04$	$3.98{\pm}0.03$

in estimating the posterior distributions, which is clearly reflected by the smaller variance of the posterior samples in Table 1.

**Tuning for quantum control** Typically, the reliable capability to manipulate qbit states is critical to quantum technologies. For instance, radio-frequency pulses can be used to change of state of spin-up and spin-down, and patients benefit from MRI scans that use these approaches. In this example, we use the BED to simulate a tuning process such that we can control the desired duration and frequency of pulses to flip electron spins in a reliable scheme. For the implicit model without access to gradients, we thus compare our proposed SAGABED method with the Bayesian optimization method. As shown in Figure 2, our method outperforms the Bayesian optimization in terms of better measurement designs in Figure 2(a) and a much smaller variance in the posterior distributions of the high-dimensional cases such as N = 50, 100 and 500 in Figure 2 (b).



Figure 2: Performance comparison between SAGABED (a-b) and BO (c-d) method on BED for tuning quantum pulse example. The contour images ((a) and (c)) show the model photon counts for optically detected spin manipulation for pulse durations (x-axis) and amounts of detuning from the spin's natural resonance frequency. (b) and (d) display the evolution of the posterior distributions with the number of designed measurements. The red points in (b) and (d) are the true mean photon count.

**Related works**. Foster et al. [4] recently proposed to use the MI lower bound for Bayesian optimal experimental design. This study relies on variational approximations to the likelihood and posterior but it is a two-stage approach where the optimal designs were determined by a separate BO. As a result, this approach has a limitation in scaling to high-dimensional design problems. A follow-up study developed by [5] aims to address the scalability issue by introducing a unified stochastic gradient-based approach. However, they assumed the models with the tractable likelihood or the gradient approximations are available. In the scope of BED for implicit models, Ao and Li [19] proposed an approximate KLD based BED method for models with intractable likelihoods; Kleinegesse and Gutmann [6, 20] have recently considered the use of the MI combined with likelihood-free inference by ratio estimation to approximate posterior distributions but this method is often computationally intensive. The authors rectify this in a follow-up study [1] that leverages MINE to jointly determine the optimal design and the posterior. However, these methods relied on a outer-level optimizer such as BO are difficult to address the scalability issues.

# 4 Conclusion

In this paper, we develop a hybrid gradient method that leverages recent developments in variational MI estimator and evolution strategies. The proposed SAGABED method incorporating a smoothed MI estimator with a guided ES achieves a unified scalable procedure to simultaneously determine the optimal design and NN parameters for implicit models without a sampling path gradient. The performance is demonstrated by one classical noisy linear model and one scientific quantum control example. The results show that our proposed method outperforms the other two baselines in terms of the MI lower bound estimate and the variance of the posterior samples given the optimal design.

### **5** Acknowledgments

This work was supported by the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research (ASCR), Applied Mathematics program under contract ERKJ352, ERKJ369; and by the Artificial Intelligence Initiative at the Oak Ridge National Laboratory (ORNL). ORNL is operated by UT-Battelle, LLC., for the U.S. Department of Energy under Contract DEAC05-00OR22725.

# 6 Broader impact

Since this work belongs to development of fundamental machine learning algorithms, it does not present any immediately foreseeable societal consequence. However, further development of the proposed method may lead to some potential positive or negative impacts on the society. Some positive impacts include (1) enable scientists and engineers to use their implicit simulator code to solve high dimensional Bayesian experiment design for scientific discovery; (2) help experimental scientist design more efficient experiments to reduce experiment cost. Furthermore, we should be cautious of the consequence of failure of the method which could cause incorrect scientific predictions that may delay the progress of scientific discovery.

### References

- [1] Steven Kleinegesse and Michael U Gutmann. Bayesian experimental design for implicit models by mutual information neural estimation. *arXiv preprint arXiv:2002.08129*, ICML 2020.
- [2] Kathryn Chaloner and Isabella Verdinelli. Bayesian experimental design: A review. *Statistical Science*, pages 273–304, 1995.
- [3] Paola Sebastiani and Henry P Wynn. Maximum entropy sampling and optimal bayesian experimental design. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 62(1):145–157, 2000.
- [4] Adam Foster, Martin Jankowiak, Elias Bingham, Paul Horsfall, Yee Whye Teh, Thomas Rainforth, and Noah Goodman. Variational bayesian optimal experimental design. In *Advances in Neural Information Processing Systems*, pages 14036–14047, 2019.
- [5] Adam Foster, Martin Jankowiak, Matthew O'Meara, Yee Whye Teh, and Tom Rainforth. A unified stochastic gradient approach to designing bayesian-optimal experiments. In *International Conference on Artificial Intelligence and Statistics*, pages 2959–2969. PMLR, 2020.
- [6] Steven Kleinegesse and Michael U Gutmann. Efficient bayesian experimental design for implicit models. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 476–485, 2019.
- [7] Antony Overstall, James McGree, et al. Bayesian design of experiments for intractable likelihood models using coupled auxiliary models and multivariate emulation. *Bayesian Analysis*, 2018.
- [8] Sophie Harbisher, Colin S Gillespie, and Dennis Prangle. Bayesian optimal design using stochastic gradient optimisation and fisher information gain. *arXiv preprint arXiv:1904.05703*, 2019.
- [9] Mohamed Ishmael Belghazi, Aristide Baratin, Sai Rajeswar, Sherjil Ozair, Yoshua Bengio, Aaron Courville, and R Devon Hjelm. Mine: mutual information neural estimation. *arXiv* preprint arXiv:1801.04062, 2018.
- [10] Sebastian Nowozin, Botond Cseke, and Ryota Tomioka. f-gan: Training generative neural samplers using variational divergence minimization. In *Advances in neural information processing systems*, pages 271–279, 2016.
- [11] Ben Poole, Sherjil Ozair, Aaron van den Oord, Alexander A Alemi, and George Tucker. On variational bounds of mutual information. *arXiv preprint arXiv:1905.06922*, 2019.

- [12] Jiaming Song and Stefano Ermon. Understanding the limitations of variational mutual information estimators. arXiv preprint arXiv:1910.06222, 2019.
- [13] Tim Salimans, Jonathan Ho, Xi Chen, Szymon Sidor, and Ilya Sutskever. Evolution strategies as a scalable alternative to reinforcement learning. *arXiv preprint arXiv:1703.03864*, 2017.
- [14] Krzysztof M Choromanski, Aldo Pacchiano, Jack Parker-Holder, Yunhao Tang, and Vikas Sindhwani. From complexity to simplicity: Adaptive es-active subspaces for blackbox optimization. In Advances in Neural Information Processing Systems, pages 10299–10309, 2019.
- [15] Krzysztof Choromanski, Mark Rowland, Vikas Sindhwani, Richard E Turner, and Adrian Weller. Structured evolution with compact architectures for scalable policy optimization. arXiv preprint arXiv:1804.02395, 2018.
- [16] Jiaxin Zhang, Hoang Tran, Dan Lu, and Guannan Zhang. A novel evolution strategy with directional gaussian smoothing for blackbox optimization. *arXiv preprint arXiv:2002.03001*, 2020.
- [17] Fei-Yu Liu, Zi-Niu Li, and Chao Qian. Self-guided evolution strategies with historical estimated gradients.
- [18] Niru Maheswaranathan, Luke Metz, George Tucker, Dami Choi, and Jascha Sohl-Dickstein. Guided evolutionary strategies: escaping the curse of dimensionality in random search. 2018.
- [19] Ziqiao Ao and Jinglai Li. An approximate {KLD} based experimental design for models with intractable likelihoods. *arXiv preprint arXiv:2004.00715*, AISTATS 2020.
- [20] Steven Kleinegesse, Christopher Drovandi, and Michael U Gutmann. Sequential bayesian experimental design for implicit models via mutual information. arXiv preprint arXiv:2003.09379, 2020.